GreMuTRRR: A Novel Genetic Algorithm to Solve Distance Geometry Problem for Protein Structures

Md. Lisul Islam[1,2], Swakkhar Shatabda[1] and M Sohel Rahman[2]
[1]Department of Computer Science and Engineering, United International University
House 80, Road 8A, Dhanmondi, Dhaka-1209, Bangladesh
[2]Department of Computer Science and Engineering,
Bangladesh University of Engineering and Technology, Palashi, Dhaka-1000, Bangladesh
Email: {lisul,swakkhar}@cse.uiu.ac.bd, msrahman@cse.buet.ac.bd

arXiv:1411.4246v1 [cs.NE] 16 Nov 2014

Abstract

Nuclear Magnetic Resonance (NMR) Spectroscopy is a widely used technique to predict the native structure of proteins. However, NMR machines are only able to report approximate and partial distances between pair of atoms. To build the protein structure one has to solve the Euclidean distance geometry problem given the incomplete interval distance data produced by NMR machines. In this paper, we propose a new genetic algorithm for solving the Euclidean distance geometry problem for protein structure prediction given sparse NMR data. Our genetic algorithm uses a greedy mutation operator to intensify the search, a twin removal technique for diversification in the population and a random restart method to recover stagnation. On a standard set of benchmark dataset, our algorithm significantly outperforms standard genetic algorithms.

keywords- Genetic Algorithms, Protein Structures, NMR, Euclidean Distance Geometry, Greedy Mutation

1 Introduction

The function of a protein depends on its *native conformation* or the stable three dimensional structure with minimum free energy in a particular environment. Knowledge about this native structure is of paramount importance and can have an enormous impact on the field of rational drug discovery. Nuclear magnetic resonance (NMR) spectroscopy is one of the most widely used techniques to predict native structure of proteins. NMR machines provide inter-atomic distance data for a given protein. Thus, the molecular distance geometry problem (MDGP) arises in this context: given the set of Euclidean distances between the atoms in a protein, MDGP asks to find the Cartesian coordinates of the atoms.

However, in practice, NMR machines are able to produce the inter-atomic distances of only a subset of the pairs of atoms that are spatially close and this data too lacks accuracy. As a result, we are given the upper and lower bounds of only a subset of the Euclidean distances. Thus we are left to solve a variant of the MDGP problem with incomplete and inaccurate data. Many computational approaches have been applied to solve

MDGP problem with sparse and inaccurate data on real instances [1, 2, 3, 4, 5, 6]. However, complete search methods like spatial branch and bound (sBB) and stochastic methods like variable neighborhood search (VNS) have been able to solve the problem for proteins with only upto 50 atoms [7] and fail to converge quickly for larger number of atoms.

In this paper, we present GreMuTRRR (pronounced grey matter), which is a scalable genetic algorithm to solve the MDGP problem for sparse and inaccurate NMR data. Much of the success of our approach comes from a greedy mutation operator, sude to intensify the search, a twin removal technique to diversify the population and a random restart method to recover from stagnation. Experimental results on proteins with number of atoms ranging from 50–2147 shows that our algorithm outperforms standard genetic algorithms and thus obtains state-of-the-art results.

2 Preliminaries

In this section, we provide a formal definition of the MDGP problem and a brief description of genetic algorithms.

2.1 Distance Geometry Problem

Molecular distance geometry problem can formally be defined as followed [8]: find a set of Cartesian coordinates $c_1, c_2, \cdots, c_n \in \mathbb{R}^3$ of atoms of a molecule such that $l_{ij} \leq d_{ij} \leq u_{ij}, \forall (i,j) \in E$, where l_{ij} and u_{ij} are the lower and upper bounds of the Euclidean distance $d_{ij} \equiv ||c_i - c_j||$ between a pair of atoms $(i,j) \in E$. Notably, in the context of NMR data, E is sparse. Here, Cartesian coordinate c_i of an atom i corresponds to a three dimensional point (x_i, y_i, z_i) in the Euclidean space.

2.2 Genetic Algorithm

Genetic Algorithms are population-based search methods that resemble the natural phenomena of biological evolution. Genetic Algorithms are widely used for different search optimization problems in different fraternity. It basically starts with a pool of initial random solutions which is called the initial population. Each individual in the population are encoded by a set of properties which are called chromosome or genotype which can be altered for attaining diversification in the population. It then follows an iterative process in which each of these iteration is called a generation. In each generation, the population are then allowed to evolve using different operators which also mimics the natural process of biological evolution like mutation, recombination or survival of the fittest. In each generation, the fitness of each individual is evaluated. Generally the fitness is the value of the optimization function being considered to be solved. The more fit individuals are usually selected to breed and generate new fitter individuals in the population. Thus a new generation of population is 'breeded' and are used in the next generation. This process of evolution are continues until a sufficient number of generations has

been produced or a satisfactory level of fitness value has been attained.

3 Related Works

Euclidean distance geometry problem and its variants are applied to many problems in various fields including wireless sensor network localization [9], inverse kinematic problem [10], multi dimensional scaling [11], protein structure determination [6] etc. The variant of the MDGP problem, when we know distances for all pairs $(i, j) \in E = \{1, 2, \cdots\}^2$ and $d_{ij} = l_{ij} = u_{ij}$ has a polynomial algorithm to produce exact solution [8]. Even when some of the pairwise distances are unknown, the problem is solvable by a linear time algorithm [3]. However, the variant of MDGP with sparse and inaccurate data is shown to be NP-hard by More and Wu [12]. A recent survey of computational methods applied to solve this variant of MDGP can be found in [2].

Among the general purpose methods, spatial branch and bound [13] and variable neighborhood search (VNS) [14] methods are not scalable [7]. Smoothing based methods like DGSOL [6, 12] also fail for large instances of the problem. In [15], VNS was combined with DGSOL approaches which provided better results for larger instances but resulted into a slow algorithm. A combinatorial build up algorithm was proposed in [16]. It is important to note that all these methods were tested on dense instances only. Among other methods applied to this problem graph decomposition methods [1] and NLP formulations [17] are notable.

4 Our Method

GreMuTRRR is formally presented in Algorithm 1. It starts with a population initialized randomly and terminates at convergence. Convergence is achieved if the search does not improve the quality of the global solutions for a given period of time. In each generation, individuals are selected from the population for cross-over to produce new individuals. Mutation operations are performed probabilistically on the newly found individuals. Global best solutions are updated in each iteration and are kept in the new populations to maintain elitism. Periodical twin removal and random restart operations are activated to ensure diversification and recover stagnation. Rest of this section describes the various components of GreMuTRRR.

4.1 Search Space

A Protein structure is formed by the interaction among different amino acids tied with each other by chemical bonds. A pair of these amino acid cannot reside closer than 3.8Å to each other to avoid steric clash. If all the amino acids in a certain protein structure align in a straight line, we have a chain of amino acids of length $3.8 \times V$ in each directions, where V is the number of amino acids in the protein structure. Hence, we have taken the upper bound of our 3-D search space to $3.8 \times V$.

Algorithm 1: GreMuTRRR()

```
1  nonDiverseSteps = 0
2  nonImprovingSteps = 0
3  GreedyMutationRate = 0.8
4  Intialize the population, P randomly
5  while Termination criterion are met do
6      P_new = {globalBest}
7      for each Individual X ∈ P do
8          ⟨X_1, X_2⟩ = tournamentSelection(P)
9          X_new = crossOver(X_1, X_2)
10         add X_new to P_new
11     end
12     for each Individual X in P_new do
13         if rand(0, 1) ≤ GreedyMutationRate then
14             greedyMutate(X)
15         else
16             randomMutate(X)
17         end
18     end
19     find the individual X_best ∈ P_new with best fitness
20     if fitness(globalBest) < fitness(X_best) then
21         globalBest = X_best
22         nonImprovingSteps = 0
23     else
24         nonImprovingSteps + +
25     end
26     if nonDiverseSteps ≥ twinRemovalInterval then
27         activate twinRemoval(P_new) procedure
28         nonDiverseSteps = 0
29     else
30         nonDiverseSteps + +
31     end
32     if nonImprovingSteps ≥ randomRestartInterval then
33         activate randomRestart(P_new) procedure
34         nonImprovingSteps = 0
35     end
36     P = P_new
37 end
38 return globalBest
```

4.2 Encoding

In our method, we have encoded every individual in a population by $3 \times V$ number of genes. Here, V is the number of atoms in a protein structure. Each of these genes in an individual is initialized randomly from an uniform distribution of the range $[0, 3.8 \times V]$. So, an individual X can be represented as the ordered list of $3V$ number of genes:

$$X = \left\{ \underbrace{x_1, y_1, z_1}_{1^{st}Atom}, \underbrace{x_2, y_2, z_2}_{2^{nd}Atom}, \ldots, \underbrace{x_i, y_i, z_i}_{i^{th}Atom}, \ldots, \underbrace{x_V, y_V, z_V}_{V^{th}Atom} \right\}$$

4.3 Fitness Evaluation

We have calculated the euclidean distance between each pair of points present in the individual's chromosome where each non-overlapping consecutive subsequence of length three in the chromosome represents the position of an atom in the protein structure in a 3-D space. We have assumed an upper bound, u_{ij} and lower bound, l_{ij} of the distances between pair of amino acids. The fitness of an individual X is defined by as Equation 1 below:

$$Fitness(X) = \left(\frac{1}{|E|} \sum_{(i,j) \in E} e_{ij}^2 \right)^{1/2} \tag{1}$$

where, $e_{ij} = max\{l_{ij} - \|c_i - c_j\|, \|c_i - c_j\| - u_{ij}, 0\}$ is the error associated to the constraints $l_{ij} \leq \|c_i - c_j\| \leq u_{ij}$ and $|E|$ denotes the number of distance pairs given. This is known as the *Largest Distance Error (LDE)* in the literature [2].

4.4 Genetic Operators

Genetic algorithms are guided by different genetic operators to help in keeping a balance of the exploitation and exploration of the search process. We use three types of genetic operators in our search: random mutation, greedy mutation and cross-over.

4.4.1 Random Mutation

Genetic diversity is a necessity for the process of evolution which is generally attained by mutation operator. Mutation helps the evolutionary process to guide through and look for different avenues and solutions in the search space. It also helps to avoid local optima by preventing the individuals from becoming too similar to one another. We have used uniform mutation in our method. We have mutated the value of each gene into a new value within the pre-specified bound of the search space($[0, 3.8 \times V]$). Whether an individual will be mutated or not depends on a mutation rate which is kept to a lower value of 0.015 to avoid primitive random search. A sketch of the pseudo-code is given in Algorithm 2.

Algorithm 2: Random Mutation (individual X)

```
1 mutationRate = 0.015
2 for each gene X(i) in the genotype of X do
3     if rand(0, 1) ≤ mutationRate then
4         X(i) = U(0, 1) ∗ (3 ∗ V);
5     end
6 end
7 return X
```

4.4.2 Greedy Mutation

In our proposed method, we have also used a greedy mutation where we greedily choose the new value of a particular gene in an individual. We randomly alter a certain gene and try r different random values for the gene temporarily. The value of gene that gives the best fitness of that individual is retained and used as the final new value of the gene. As, greedy mutation is computationally expensive, we make selection between the random mutation and greedy mutation with a probability, $GreedyMutationRate(= 0.8)$. The algorithm for greedy mutation is given in Algorithm 3.

Algorithm 3: Greedy Mutation (individual X)

```
1 set r = 20
2 for each gene X(i) in the genotype of X do
3     S = set of r random values for gene X(i) from the range [0, 3.8 ∗ V]
4     find v ∈ S for which fitness(X) is minimum
5     set, X(i) = v
6 end
7 return X
```

4.4.3 Cross-over

Cross-over operators help genetic process to exploit the better solutions found thus far along the evolution process and regenerate new individuals by combining the genetic information of the individuals with better fitness value. It uses the better historical genetic information to guide the search into the search space regions with solutions having better fitness. By recombining the fitter individuals to generate new offspring, this operator is likely to produce more fitter individuals. We have used tournament selection with tournament size being equal to 5 to select two individual parents to take part in the cross-over operation. We then have recombined genes of the two parents and have generated a new set of genes for each of the offsprings with equal probability for each gene of being selected from any of the parents. The algorithm is outlined in Algorithm 4.

Algorithm 4: Cross-Over (individual X_1, individual X_2)

X_{new} be a new offspring
$uniformRate = 0.5$
for *each gene X(i) in the genotype of* X_{new} **do**
    if $rand(0,1) \geq uniformRate$ **then**
        $X_{new}(i) = X_1(i)$
    else
        $X_{new}(i) = X_2(i)$
    end
end
return X_{new}

4.5 Twin Removal

We have removed and reinitialized individuals with identical genetic information. We have defined the similarity measure between two individuals as follows:

$$Similarity(X_1, X_2) = e^{-\frac{\|X_1 - X_2\|^2}{2\sigma^2}} \tag{2}$$

Here, the value of σ has been chosen to be 75. The $Similarity(X_1, X_2)$ function will return a value in the range [0,1]. The more the value of the similarity function is closer to 1, the more genetically similar are X_1 and X_2. We have chosen 0.8 as the threshold to define similarity between two individual. If $Similarity(X_1, X_2) >= 0.8$, we declare that X_1 and X_2 are twins and reinitialize randomly one of them. We have run the this Twin Removal procedure after every 100 generations of evolution.

Algorithm 5: Twin Removal

$similarityThreshold = 0.8$
for *each pair of individuals* (X_i, X_j) *in the population* **do**
    if $Similarity(X_i, X_j) \geq similarityThreshold$ **then**
        declare X_i and X_j as Twins
        reinitialize X_j
    end
end

4.6 Random Restart

If the algorithm does not show enough improvement within a significant number of generations, we reinitialize some individuals of the population with random values within the search space. We have evaluated the improvement over the immediate past 50 generations. If the improvement in the last 50 generations is less than or equal to a threshold, $t = 0.001$, we initiate the random restart procedure.

5 Experimental Results

We have implemented GreMuTRRR in Java programming language using JDK 1.6 and have run our experiments on an Intel 3.3GHz core $i3$ machine with 2GB RAM. We compared the performance of GreMuTRRR with a basic genetic algorithm without considering the greedy mutation, twin removal and random restart. We denote this algorithm as 'Basic GA' henceforth. Basic GA differs with GreMuTRRR given in Algorithm 1 in Lines 27–36 since it does not contain twin removal and random restart and in Line 3, as the parameter $greedyMutationRate = 0.0$ for Basic GA. All other parameters were kept same for the sake of fair comparison.

5.1 Benchmark

We have considered two sets of protein instances for our experiments: one considering only the backbone atoms (backbone only) and another containing all atoms (full atomic). These are larger protein benchmarks introduced in [18]. We extracted the structures from PDB website [19] and have calculated the distance among pair of atoms. Lower and upper bounds of the distances are calculated using the following equation:

$$l_{ij} = (1 - \epsilon)\hat{d}_{ij}, u_{ij} = (1 + \epsilon)\hat{d}_{ij} \tag{3}$$

Here, $\hat{d}_{ij}$ is the real distance between point c_i and point c_j in the known structure of the protein sequence and $\epsilon = 0.8$. To make the dataset sparse, we took only 30% of the distances that are less or equal to $6\mathring{A}$.

5.2 Results

The largest distance error (LDE) values for our algorithm and Basic GA for backbone only instances and full atomic instances are given in Table 1 and Table 2 respectively. Each protein is reported with its corresponding PDB id and number of atoms in the structure considered.

Table 1: Best Fitness for population size 50 with run of 2000 generations(with backbone atoms only)

Protein Id	V	Basic GA	GreMuTRRR
1PTQ	50	0.49546	0.44977
1LFB	77	0.18060	0.13360
1F39	101	0.06595	0.02051
1AX8	130	0.06064	0.00253
1RGS	264	0.12380	0.00297
1TOA	277	0.12993	0.00156
1KDH	356	0.19119	0.00458
1BPM	481	0.34434	0.00441
1MQQ	679	0.58863	0.00894

Table 2: Best Fitness for population size 50 with run of 1500 generations(including all atoms)

Protein Id	V	Basic GA	GreMuTRRR
1PTQ	402	0.71000	0.17224
1LFB	641	0.97121	0.31463
1F39	767	1.81091	0.16805
1AX8	1003	2.52303	0.43211

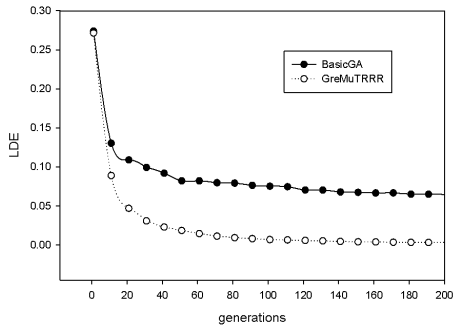

Figure 1: Search progress for two algorithms for the protein 1AX8.

5.3 Analysis

From the results presented in the tables it is evident that for both of the instances GreMuTRRR achieves better solutions compared to the Basic GA algorithm for the proteins. However, the improvement is more clearly shown in Figure 1. GreMuTRRR significantly lowers the objective function LDE before both algorithms gets stuck and can not improve further. Its interesting to note thatGreMuTRRR is able to converge very early (≤ 50 generation) compared to the Basic GA.

6 Conclusion

In this paper, we presented a new genetic algorithm to solve the molecular distance geometry for protein structure determination problem using NMR data. We use a greedy mutation operator to intensify the search, a twin removal technique for diversification in the population and a random restart method to recover stagnation. On a standard set of benchmark dataset, our algorithm significantly outperforms standard genetic algorithms. In future, we want to implement a web service based on our method and make it public to be used by biologists.

References

[1] M. Souza, C. Lavor, A. Muritiba, and N. Maculan, "Solving the molecular distance geometry problem with inaccurate distance data," *BMC bioinformatics*, vol. 14, no. Suppl 9, p. S7, 2013.

[2] L. Liberti, C. Lavor, N. Maculan, and A. Mucherino, "Euclidean distance geometry and applications," *SIAM Review*, vol. 56, no. 1, pp. 3–69, 2014.

[3] D. Wu and Z. Wu, "An updated geometric build-up algorithm for solving the molecular distance geometry problems with sparse distance data," *Journal of Global Optimization*, vol. 37, no. 4, pp. 661–673, 2007.

[4] A. Mucherino, L. Liberti, C. Lavor, and N. Maculan, "Comparisons between an exact and a metaheuristic algorithm for the molecular distance geometry problem," in *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*. ACM, 2009, pp. 333–340.

[5] L. Liberti, C. Lavor, A. Mucherino, and N. Maculan, "Molecular distance geometry methods: from continuous to discrete," *International Transactions in Operational Research*, vol. 18, no. 1, pp. 33–51, 2011.

[6] J. J. Moré and Z. Wu, "Distance geometry optimization for protein structures," *Journal of Global Optimization*, vol. 15, no. 3, pp. 219–234, 1999.

[7] C. Lavor, L. Liberti, and N. Maculan, "Computational experience with the molecular distance geometry problem," in *Global Optimization*. Springer, 2006, pp. 213–225.

[8] G. M. Crippen and T. F. Havel, *Distance geometry and molecular conformation*. Research Studies Press Taunton, UK, 1988, vol. 74.

[9] C. Savarese, J. M. Rabaey, and J. Beutel, "Location in distributed ad-hoc wireless sensor networks," in *Acoustics, Speech, and Signal Processing, 2001. Proceedings.(ICASSP'01). 2001 IEEE International Conference on*, vol. 4. IEEE, 2001, pp. 2037–2040.

[10] D. Tolani, A. Goswami, and N. I. Badler, "Real-time inverse kinematics techniques for anthropomorphic limbs," *Graphical models*, vol. 62, no. 5, pp. 353–388, 2000.

[11] J. B. Tenenbaum, V. De Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.

[12] J. J. Moré and Z. Wu, "Global continuation for distance geometry problems," *SIAM Journal on Optimization*, vol. 7, no. 3, pp. 814–836, 1997.

[13] L. Liberti and S. Kucherenko, "Comparison of deterministic and stochastic approaches to global optimization," *International Transactions in Operational Research*, vol. 12, no. 3, pp. 263–285, 2005.

[14] L. Liberti and M. Drazic, "Variable neighbourhood search for the global optimization of constrained nlps," in *Proceedings of GO Workshop, Almeria, Spain*, vol. 2005, 2005.

[15] L. Liberti, C. Lavor, N. Maculan, and F. Marinelli, "Double variable neighbourhood search with smoothing for the molecular distance geometry problem," *Journal of Global Optimization*, vol. 43, no. 2-3, pp. 207–218, 2009.

[16] Q. Dong and Z. Wu, "A geometric build-up algorithm for solving the molecular distance geometry problem with sparse distance data," *Journal of Global Optimization*, vol. 26, no. 3, pp. 321–333, 2003.

[17] B. Hendrickson, "The molecule problem: Exploiting structure in global optimization," *SIAM Journal on Optimization*, vol. 5, no. 4, pp. 835–857, 1995.

[18] P. Biswas, K.-C. Toh, and Y. Ye, "A distributed sdp approach for large-scale noisy anchor-free graph realization with applications to molecular conformation," *SIAM Journal on Scientific Computing*, vol. 30, no. 3, pp. 1251–1277, 2008.

[19] H. M. Berman, J. D. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov, and P. E. Bourne, "The protein data bank," *Nucleic Acids Research*, vol. 28, no. 1, pp. 235–242, 2000.